

SYSTEM AND METHOD FOR PROVIDING LAST-
MILE DATA PRIORITIZATION

BACKGROUND OF THE INVENTION

1. Related Applications.

5 The present invention claims priority from U.S. Provisional Patent Application No. 60/197,490 entitled CONDUCTOR GATEWAY filed on April 17, 2000.

2. Field of the Invention.

10 The present invention relates, in general, to network information access and, more particularly, to software, systems and methods for serving web pages in a coordinated fashion from multiple cooperating web servers.

3. Relevant Background.

15 Increasingly, business data processing systems, entertainment systems, and personal communications systems are implemented by computers across networks that are interconnected by internetworks (e.g., the Internet). The Internet is rapidly emerging as the preferred system for distributing and exchanging data. Data exchanges support
20 applications including electronic commerce (e-commerce), broadcast and multicast messaging, videoconferencing, gaming, and the like.

25 The Internet is a collection of disparate computers and networks coupled together by a web of interconnections using standardized communications protocols. The Internet is characterized by its vast reach as a result of its wide

and increasing availability and easy access protocols. Unfortunately, the ubiquitous nature of the Internet results in variable bandwidth and quality of service between points. The latency and reliability of data transport is largely determined by the total amount of traffic on the Internet and so varies wildly seasonally and throughout the day. Other factors that affect quality of service include equipment outages and line degradation that force packets to be rerouted, damaged and/or dropped. Also, routing software and hardware limitations within the Internet infrastructure may create bandwidth bottlenecks even when the mechanisms are operating within specifications.

Users connect to the Internet through a variety of mechanisms, but many mechanisms use shared bandwidth connections. For example, data connections using broadband wireless access (BWA), cellular telephones and cable modems typically share a given amount of bandwidth amongst a plurality of users. Each user's data is transported on a first-in first-out (FIFO) basis. In such a transport medium, there is no mechanism for prioritizing some data over other data, or some users over other users. Prioritization has not been an issue in many applications because the average latency of such networks has been sufficiently low and sufficiently uniform to provide acceptable performance. However, there is an increasing demand for network applications and data access that cannot tolerate high and variable latency.

In e-commerce applications, it is important to provide a satisfying user experience that leads to a purchase transaction. To provide this high level of service, a web site operator must ensure that data is delivered to the customer in the most usable, efficient,

and timely fashion. Also, the web site operator must ensure that critical data received from the customer is handled with priority. Until now, however, the e-commerce site owner has had little or no control over the transport mechanisms from the Internet to a user device that affect the latency and quality of service.

While efforts are continually being made to increase the capacity and quality of service afforded by the Internet, it is contemplated that congestion will always impact the ability to predictably and reliably offer a specified level of service. This is particularly true in wireless and other shared bandwidth transport media that historically oversubscribe the available bandwidth. Moreover, the change in the demand for bandwidth increases at a greater rate than does the change in bandwidth supply, ensuring that congestion will continue to be an issue into the foreseeable future. A need exists for a system to exchange data that provides a high quality of service even during periods of congestion.

The communication mechanisms linking a user to the Internet are often referred to as the "last-mile". Because the last-mile connectivity must fan out to reach a large number of users, there is great emphasis on shared bandwidth mechanisms that enable any infrastructure investment to be leveraged over many users. The last-mile provider (e.g., a broadband wireless access provider, cable company, or cellular telephone company) provides routers that couple to the Internet, and wireless transceivers that couple between the routers and users to implement the shared bandwidth medium.

In general, once data communication leaves the confines of the Internet into a shared bandwidth medium, web sites and infrastructure entities that attempt to

provide prioritization loose the ability regulate and control the flow of data communication over the last-mile. While the last-mile provider may attempt to provide bandwidth management within their network, there currently is no mechanism for an Internet entity, such as a web site, to readily manipulate the shared bandwidth medium.

SUMMARY OF THE INVENTION

Briefly stated, the present invention involves systems and methods for prioritizing data traffic over a shared bandwidth connection. In one embodiment, traffic is prioritized implicitly by governing the rate at which data is launched into the shared bandwidth connection. In another embodiment, mechanisms that couple the Internet to the shared bandwidth connection are explicitly regulated through a control interface. Priority information is coupled through the Internet to front-end computers. The front-end computers couple to the shared bandwidth connection and implement the implicit and/or explicit regulation.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates a general distributed computing environment in which the present invention is implemented;

FIG. 2 shows in block-diagram form significant components of a system in accordance with the present invention;

FIG. 3 shows front-end components of FIG. 2 in greater detail; and

FIG. 4 shows the data blender prioritization mechanism of FIG. 3 in greater detail.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

The present invention is directed to systems and methods for prioritizing data traffic as it is communicated between a network such as the Internet, and an end user's client computer. This communication link is often referred to as the "last-mile" in network architecture descriptions. The last-mile link is typically implemented using different communication technology than is used in the Internet itself. For example, modem connections to a modem pool maintained by an Internet service provider (ISP) were prevalent for many years. More recently, last-mile links are being implemented using broadband communication technologies such as digital subscriber line (DSL), broadband wireless, cable modems, as well as digital cellular communication.

These technologies characteristically share available bandwidth amongst a plurality of end-user machines in a manner controlled by the last-mile provider (e.g., the ISP, cellular system operator, cable company and the like). The present invention is directed to systems and methods enabling entities outside of the last-mile provider's system to have some level of control over the effective allocation of bandwidth within the last-mile system. In some embodiments, the effective bandwidth allocation is achieved without any presence within the last-mile communication system itself.

The present invention is illustrated and described in terms of a distributed computing environment such as an enterprise computing system using public communication channels such as the Internet. However, an important feature of the present invention is that it is readily scaled upwardly and downwardly to meet the needs of a particular application. Accordingly, unless specified to

the contrary, the present invention is applicable to significantly larger, more complex network environments, including wireless network environments, as well as small network environments such as conventional and wireless LAN systems.

In accordance with the disclosed embodiments of the present invention, prioritization performed by a front-end server located at the edge of the network. The front-end server is coupled to communicate data, typically in the form of IP packets, with an interface to the last-mile communication system, although any communication protocol supported by the last-mile interface can be adopted by the front-end server. The front-end server implements a mechanism for discriminating between request and/or response packets communicated with the last-mile communication system. In the case of outgoing traffic (i.e., originating from the last-mile interface), the front-end preferably buffers and/or reorders requests before transmitting them over the Internet. In the case of incoming traffic (i.e., originating from the Internet), the front-end server implements a mechanism for buffering and/or reordering response packets before they are transmitted over the last-mile interface.

The front-end server either explicitly or implicitly controls the mix of data flowing through the interface to the last-mile communication system. In one embodiment, the front-end is coupled to supply data to the last-mile interface, and can therefore implicitly prioritize data by regulating the order and rate at which data is supplied to the interface. In another embodiment, the front-end establishes policy for a router that feeds data to the last-mile interface. This enables the front-end to specifically manipulate the order and rate at which

packets pass through the router and are applied to the last-mile interface.

A particular last-mile communication system may be statically associated with a front-end server so that all communication entering and leaving the last-mile interface goes through the associated front-end server. Alternatively, a redirection mechanism is enabled to select from a pool of available front-end servers and client request-response traffic is conducted through the selected front-end server.

The front-end server 201 handles communication with originating servers or other network servers that contain the information and resources that are the subject of each request. In some embodiments, it is useful for the front-end server 201 to indicate priority information associated with particular data so that these other network servers may use the priority information. For example, resource reservation protocol (RSVP) and integrated services described in IETF document RFC 1633 are mechanisms that provide various levels of service based upon an indicated priority. This enables the client 205 to benefit from prioritization functionality implemented elsewhere in the data network without necessarily requiring modification to the client 205.

Preferably, the front-end server establishes and maintains an enhanced communication channel with the originating web servers or other network resources. By enhanced, it is meant that the channel offers improved quality of service, lower latency, prioritization services, higher security transport, or other features and services that improve upon the basic transport mechanisms (such as TCP) defined for Internet data transport. In this manner, prioritization functionality can be provided

both in the Internet, and in the last-mile connection to the Internet.

For purposes of this document, a web server is a computer running server software coupled to the World Wide Web (i.e., "the web") that delivers or serves web pages. The web server has a unique IP address and accepts connections in order to service requests by sending back responses. A web server differs from a proxy server or a gateway server in that a web server has resident a set of resources (i.e., software programs, data storage capacity, and/or hardware) that enable it to execute programs to provide an extensible range of functionality such as generating web pages, accessing remote network resources, analyzing contents of packets, reformatting request/response traffic and the like using the resident resources. In contrast, a proxy simply forwards request/response traffic on behalf of a client to resources that reside elsewhere, or obtains resources from a local cache if implemented. A web server in accordance with the present invention may reference external resources of the same or different type as the services requested by a user, and reformat and augment what is provided by the external resources in its response to the user. Commercially available web server software includes Microsoft Internet Information Server (IIS), Netscape Netsite, Apache, among others. Alternatively, a web site may be implemented with custom or semi-custom software that supports HTTP traffic.

FIG. 1 shows an exemplary computing environment 100 in which the present invention may be implemented. Environment 100 includes a plurality of local networks such as Ethernet network 102, FDDI network 103 and Token Ring network 104. Essentially, a number of computing

devices and groups of devices are interconnected through a network 101. For example, local networks 102, 103 and 104 are each coupled to network 101 through routers 109. LANs 102, 103 and 104 may be implemented using any available topology and may implement one or more server technologies including, for example UNIX, Novell, or Windows NT networks, or peer-to-peer type network. Each network will include distributed storage implemented in each device and typically includes some mass storage device coupled to or managed by a server computer. Network 101 comprises, for example, a public network such as the Internet or another network mechanism such as a fibre channel fabric or conventional WAN technologies.

Local networks 102, 103 and 104 include one or more network appliances 107. One or more network appliances 107 may be configured as an application and/or file server. Each local network 102, 103 and 104 may include a number of shared devices (not shown) such as printers, file servers, mass storage and the like. Similarly, devices 111 may be shared through network 101 to provide application and file services, directory services, printing, storage, and the like. Routers 109 provide a physical connection between the various devices through network 101. Routers 109 may implement desired access and security protocols to manage access through network 101. Network appliances 107 may also couple to network 101 through public switched telephone network 108 using copper or wireless connection technology. In a typical environment, an Internet service provider 106 supports a connection to network 101 as well as PSTN 108 connections to network appliances 107.

Network appliances 107 may be implemented as any kind of network appliance having sufficient computational

function to execute software needed to establish and use a connection to network 101. Network appliances 107 may comprise workstation and personal computer hardware executing commercial operating systems such as Unix variants, Microsoft Windows, Macintosh OS, and the like. At the same time, some appliances 107 comprise portable or handheld devices using wireless connections through a wireless access provider such as personal digital assistants and cell phones executing operating system software such as PalmOS, WindowsCE, EPOCOS, and the like. Moreover, the present invention is readily extended to network devices such as office equipment, vehicles, and personal communicators that make occasional connection through network 101.

Each of the devices shown in FIG. 1 may include memory, mass storage, and a degree of data processing capability sufficient to manage their connection to network 101. The computer program devices in accordance with the present invention are implemented in the memory of the various devices shown in FIG. 1 and enabled by the data processing capability of the devices shown in FIG. 1. In addition to local memory and storage associated with each device, it is often desirable to provide one or more locations of shared storage such as disk farm (not shown) that provides mass storage capacity beyond what an individual device can efficiently use and manage. Selected components of the present invention may be stored in or implemented in shared mass storage.

Fig. 2 illustrates a private network 200 implemented within the Internet infrastructure. Private network 200 expedites and prioritizes communications between a client 205 and a web site 210. In the specific examples herein clients 205 comprises client software that can access a

standard or proprietary user level protocol for network access. Examples include client applications for file transfer protocol (FTP) services, voice over Internet protocol (VoIP) services, network news protocol (NNTP) services, multi-purpose internet mail extensions (MIME) services, post office protocol (POP) services, simple mail transfer protocol (SMTP) services, as well as Telnet services. In addition to network protocols, the client application may access a network application such as a database management system (DBMS) in which case the client application generates query language (e.g., structured query language or "SQL") messages. In wireless appliances, a client application may communicate via a wireless application protocol (WAP), extensible markup language (XML) or the like.

For convenience, the term "web site" is used interchangeably with "web server" in the description herein although it should be understood that a web site comprises a collection of content, programs and processes implemented on one or more web servers. A web site is owned by the content provider such as an e-commerce vendor whereas a web server refers to set of programs running on one or more machines coupled to an Internet node. The web site 210 may be hosted on the site owner's own web server, or hosted on a web server owned by a third party. A web hosting center is an entity that implements one or more web sites on one or more web servers using shared hardware and software resources across the multiple web sites. In a typical web infrastructure, there are many web browsers, each of which has a TCP connection to the web server in which a particular web site is implemented. The present invention adds two components to the infrastructure: a front-end 201 and back-end 203. Front-end 201 and back-end

203 are coupled by a managed data communication link 202 that forms, in essence, a private network.

Front-end mechanisms 201 serve as an access point for client-side communications through various shared bandwidth last-mile connections. In one example, a front-end 201 couples directly to a wireless network interface 215. In another example, a front-end 201 couples directly to a cable head end unit 225 that feeds multiple clients 205 operating cable set-top boxes (not shown). In another embodiment, front-end 201 couples to a router 235 that communicates with a wireless network interface 215. In each case, front-end 201 is able to regulate the rate and/or order at which data is applied to the shared bandwidth network.

Front-ends 201 include mechanisms for discriminating between data in order to implement effective prioritization. The mechanisms can take a variety of forms, however. In one preferred implementation, response data packets arriving at the front-end 201 from the Internet are associated with a prioritization value supplied by, for example, the web site owner. The enhanced channel 202 is important in this implementation to provide a ready mechanism to communicate this prioritization value to the front-end 201. In another embodiment, front-end 201 determines a priority value based on a user identification and/or content type and/or other criteria that it can determine from a request packet received from the last-mile connection. When a corresponding response packet is received, the determined priority value is used to order the response packet relative to other users' response packets before they are applied to the last-mile connection (e.g., interface 215, head end 225 or router 235).

Transport mechanism 202 is implemented by cooperative actions of the front-end 201 and back-end 203. Back-end 203 processes and directs data communication to and from web site 210. Transport mechanism 202 communicates data packets using a proprietary protocol over the public Internet infrastructure in the particular example. Hence, the present invention does not require heavy infrastructure investments and automatically benefits from improvements implemented in the general purpose network 101. Unlike the general purpose Internet, front-end 201 and back-end 203 are programmably assigned to serve accesses to a particular web site 210 at any given time.

It is contemplated that any number of front-end and back-end mechanisms may be implemented cooperatively to support the desired level of service required by the web site owner. The present invention implements a many-to-many mapping of front-ends to back-ends. Because the front-end to back-end mappings can be dynamically changed, a fixed hardware infrastructure can be logically reconfigured to map more or fewer front-ends to more or fewer back-ends and web sites or servers as needed.

Front-end 201 together with back-end 203 function to reduce traffic across the TMP link 202 and to improve response time for selected browsers. Traffic across the TMP link 202 is reduced by compressing data and serving browser requests from cache for fast retrieval. Also, the blending of request datagrams results in fewer request:acknowledge pairs across the TMP link 202 as compared to the number required to send the packets individually between front-end 201 and back-end 203. This action reduces the overhead associated with transporting a given amount of data, although conventional request:acknowledge traffic is still performed on the

links coupling the front-end 201 to client 205 and back-end 203 to a web server. Moreover, resend traffic is significantly reduced further reducing the traffic. Response time is further improved for select privileged users and for specially marked resources by determining the priority for each HTTP transmission.

In one embodiment, front-end 201 and back-end 203 are closely coupled to the Internet backbone. This means they have high bandwidth connections, can expect fewer hops, and have more predictable packet transit time than could be expected from a general-purpose connection. Although it is preferable to have low latency connections between front-ends 201 and back-ends 203, a particular strength of the present invention is its ability to deal with latency by enabling efficient transport and traffic prioritization. Hence, in other embodiments front-end 201 and/or back-end 203 may be located farther from the Internet backbone and closer to clients 205 and/or web servers 210. Such an implementation reduces the number of hops required to reach a front-end 201 while increasing the number of hops within the TMP link 202 thereby yielding control over more of the transport path to the management mechanisms of the present invention.

Hence, clients 205 do not conduct all data transactions directly with a web server 210. Instead, clients 205 conduct some and preferably a majority of transactions with front-ends 201. Client data is then sent, using TMP link 202, to the back-end 203 and then to the web server 210. Running multiple clients 205 over one large connection provides several advantages:

- Since all client data is mixed, each client can be assigned a priority. Higher priority clients, or clients requesting higher priority data, can be

given preferential access to network resources so they receive access to the channel sooner while ensuring low-priority clients receive sufficient service to meet their needs.

- 5 • The large connection between a front-end 201 and back-end 203 can be permanently maintained, shortening the many TCP/IP connection sequences normally required for many clients connecting and disconnecting.

10 Using a proprietary protocol allows the use of more effective techniques to improve data throughput and makes better use of existing bandwidth during periods when the network is congested.

15 A particular advantage of the architecture shown in FIG. 2 is that it is readily scaled. Any number of client machines 205 may be supported. In a similar manner, a web site owner may choose to implement a site using multiple web servers 210 that are co-located or distributed throughout network 101. To avoid congestion, additional front-ends 201 may be implemented or assigned to particular web sites. Each front-end 201 is dynamically re-configurable by updating address parameters to serve particular web sites. Client traffic is dynamically directed to available front-ends 201 to provide load
25 balancing. Hence, when quality of service drops because of a large number of client accesses, an additional front-end 201 can be assigned to the web site and subsequent client requests directed to the newly assigned front-end 201 to distribute traffic across a broader base.

30 In the particular examples, this is implemented by a front-end manager component 207 that communicates with multiple front-ends 201 to provide administrative and

configuration information to front-ends 201. Each front-end 201 includes data structures for storing the configuration information, including information identifying the IP addresses of web servers 210 to which they are currently assigned. Other administrative and configuration information stored in front-end 201 may include information for prioritizing data from and to particular clients, quality of service information, and the like.

Similarly, additional back-ends 203 can be assigned to a web site to handle increased traffic. Back-end manager component 209 couples to one or more back-ends 203 to provide centralized administration and configuration service. Back-ends 203 include data structures to hold current configuration state, quality of service information and the like. In the particular examples front-end manager 207 and back-end manager 209 serve multiple web sites 210 and so are able to manipulate the number of front-ends and back-ends assigned to each web site 210 by updating this configuration information. When the congestion for the site subsides, the front-end 201 and back-end 203 can be reassigned to other, busier web sites. These and similar modifications are equivalent to the specific examples illustrated herein.

In the case of web-based environments, front-end 201 is implemented using custom or off-the-shelf web server software. Front-end 201 is readily extended to support other, non-web-based protocols, however, and may support multiple protocols for varieties of client traffic. Front-end 201 processes the data traffic it receives, regardless of the protocol of that traffic, to a form suitable for transport by TMP 202 to a back-end 203. Hence, most of the functionality implemented by front-end

201 is independent of the protocol or format of the data received from a client 205. Although the discussion of the exemplary embodiments herein relates primarily to front-end 201 implemented as a web server, it should be noted that, unless specified to the contrary, web-based traffic management and protocols are merely examples and not a limitation of the present invention.

FIG. 3 illustrates principle functional components of an exemplary front-end 201 in greater detail. Primary functions of the front-end 201 include creating and maintaining a TCP socket for each active client communication channel and prioritizing the order and rate at which data is transmitted on those sockets relative to each other. In the preferred implementations, front-end 201 also translates transmission control protocol (TCP) packets from last-mile interface 215/225/235 into transport morphing protocol™ (TMP™) packets used in the system in accordance with the preferred embodiments. Transport morphing protocol and TMP are trademarks or registered trademarks of Circadence Corporation in the United States and other countries. It is contemplated that the various functions described in reference to the specific examples may be implemented using a variety of data structures and programs operating at any location in a distributed network. For example, a front-end 201 may be operated on a network appliance 107 or server within a particular network 102, 103, or 104 shown in FIG. 1. The present invention is readily adapted to any application where multiple clients are coupling to a centralized resource. Moreover, other transport protocols may be used, including proprietary transport protocols.

TCP component 301 includes devices for implementing physical connection layer and Internet protocol (IP) layer

functionality. Current IP standards are described in IETF documents RFC0791, RFC0950, RFC0919, RFC0922, RFC792, RFC1112 that are incorporated by reference herein. For ease of description and understanding, these mechanisms are not described in great detail herein. Where protocols other than TCP/IP are used to couple to a client 205, TCP component 401 is replaced or augmented with an appropriate network protocol process.

TCP component 301 communicates TCP packets with one or more clients 205 through last-mile interface 215/225/235 as appropriate. Received packets are coupled to parser 302 where the Internet protocol (or equivalent) information is extracted. TCP is described in IETF RFC0793 which is incorporated herein by reference. Each TCP packet includes header information that indicates addressing and control variables, and a payload portion that holds the user-level data being transported by the TCP packet. The user-level data in the payload portion typically comprises a user-level network protocol datagram.

Parser 302 analyzes the payload portion of the TCP packet. In the examples herein, HTTP is employed as the user-level protocol because of its widespread use and the advantage that currently available browser software is able to readily use the HTTP protocol. Particularly in the case of wireless devices, however, other protocols may be used more commonly such as handheld device markup language (HDML), wireless access protocol (WAP) and extensible markup language (XML) protocols. Parser 302 may implement file transfer protocol (FTP), mail protocols such as simple mail transport protocol (SMTP), structured query language (SQL) and the like. Essentially, parser 302 comprises a parser that can look into the received

packet to identify information from the payload portion of the packet, regardless of the particular user-level protocol in use. Any user-level protocol, including proprietary protocols, may be implemented within the present invention using appropriate modification of parser 302.

To improve performance, front-end 201 optionally includes a caching mechanism 303. Cache 303 may be implemented as a passive cache that stores frequently and/or recently accessed web pages or as an active cache that stores network resources that are anticipated to be accessed. In non-web applications, cache 303 may be used to store any form of data representing database contents, files, program code, and other information. Upon receipt of a TCP packet, HTTP parser 302 determines if the packet is making a request for data within cache 303. If the request can be satisfied from cache 303, the data is supplied directly without reference to web server 210 (i.e., a cache hit). Cache 303 implements any of a range of management functions for maintaining fresh content. For example, cache 303 may invalidate portions of the cached content after an expiration period specified with the cached data or by web sever 210. Also, cache 303 may proactively update the cache contents even before a request is received for particularly important or frequently used data from web server 210. Cache 303 evicts information using any desired algorithm such as least recently used, least frequently used, first in/first out, or random eviction. When the requested data is not within cache 303, a request is processed to web server 210, and the returned data may be stored in cache 303.

Several types of packets will cause parser 304 to forward a request towards web server 210. For example, a

request for data that is not within cache 303 (or if optional cache 303 is not implemented) will require a reference to web server 210. Some packets will comprise data that must be supplied to web server 210 (e.g., customer credit information, form data and the like). In these instances, HTTP parser 302 couples to data blender 304.

Optionally, front-end 201 implements security processes, compression processes, encryption processes and the like to condition the received data for improved transport performance and/or provide additional functionality. These processes may be implemented within any of the functional components (e.g., data blender 304) or implemented as separate functional components within front-end 201.

Also, front-end 201 may implement prioritization logic (shown in FIG. 4) to identify packets that should be given higher priority service. A prioritization program requires only that front-end 201 include a data structure associating particular clients 205 or particular TCP packet types or contents with a prioritization value. Although the prioritization value is used primarily for determining the rate and order that response packets are sent to the last-mile communication system, the prioritization value may be put to other uses for both request and response traffic. Based on the prioritization value, front-end 201 may selectively implement such features as caching, encryption, security, compression and the like to provide differential performance and/or functionality. The prioritization value is provided by the owners of web site 210, for example, and may be dynamically altered, statically set, or updated from time to time to meet the needs of a particular application.

Blender 304 slices and/or coalesces the data portions of the received packets into a more desirable "TMP units" that are sized for transport through the TMP mechanism 202. The data portion of TCP packets may range in size depending on client 205 and any intervening links coupling client 205 to TCP component 301. Moreover, where compression is applied, the compressed data will vary in size depending on the compressibility of the data. Data blender 304 receives information from front-end manager 207 that enables selection of a preferable TMP packet size. Alternatively, a fixed TMP packet size can be set that yields desirable performance across TMP mechanism 202. Data blender 304 also marks the TMP units so that they can be re-assembled at the receiving end.

In the preferred implementation, data blender 304 also serves as a buffer for storing packets from all clients 205 that are associated with front-end 201. Blender 304 mixes data packets coming into front-end 201 into a cohesive stream of TMP packets sent to a back-end 203 over TMP link 202. In creating a TMP packet, blender 304 is able to pick and choose amongst the available data packets so as to prioritize some data packets over others.

Also shown in FIG. 3 are data filter component 306 and HTTP reassemble component 307 that process incoming (with respect to client 205) data. TMP mechanism 305 receives TMP packets from TMP pipe 202 and extracts the TMP data units. Using the appended sequencing information, the extracted data units are reassembled into HTTP data packet information by HTTP reassembler 307. Data filter component 306 may also implement data decompression where appropriate, decryption, and handle caching when the returning data is of a cacheable type.

Fig. 4 illustrates the TCP component 301 in greater detail. Each TCP socket 402 maintained between the front-end 201 and a client 205 corresponds to a location 404 within reorder buffer 403. A TCP socket is created upon receipt of a request packet from a client 205. As a response packet is received into front-end 201, it is placed in its corresponding location 404 in reorder buffer 403. Reorder buffer 403 can receive response packets from either HTTP reassemble component 307 or cache 303 depending on where the response packet is served from.

When a TCP socket is created, it is assigned a priority by prioritization logic 401. For purposes of the present invention, any algorithm or criteria may be used to assign a priority, but initially it may be assigned a default or neutral priority. Alternatively, a priority may be communicated in the payload portion of a request, or by way of a cookie value accompanying the request. The TCP socket priority may be updated by prioritization logic 401 when a response is received by way of a priority value explicitly encoded in the response or a set cookie value accompanying the response. Alternatively, the priority value may be implicitly updated by analysis of the content of the response (e.g., stock quotes are given a high priority, user input forms given a low priority). Alternatively, priority information may be provided out of band.

When sending data, prioritization logic 401 selects from among the entries 404 that are ready to be transmitted. Although front-end 201 can maintain tens or hundreds of TCP sockets 401 and preferably has a high bandwidth connection, not all ready packets can use the TCP connection at one time. Prior systems would handle this by sending packets on a first-in-first-out (FIFO)

basis and by using the fairness algorithms built into the TCP protocol. In contrast, the present invention explicitly selects certain TCP sockets on which to send data according to an independent prioritization scheme.

5 In one embodiment, prioritization is accomplished by selecting which entries 404 are sent from reorder buffer 403 to the corresponding TCP socket 402. From this point, TCP socket 402 implements standard TCP functionality to transport the packet. This functionality includes a credit-based flow control mechanism that provides a
10 fairness algorithm that associates a window of a particular size with each TCP socket. A TCP socket can only send packets when there is room in its window, and must buffer all packets until room in its window becomes
15 available. Each time a packet is sent, the available space in the window decreases and each time an acknowledge packet is received the available space in the window is increased. The window size changes by expanding gradually up to a predefined maximum size when bandwidth is
20 available, and contracting rapidly where there is congestion as indicated by lost and/or damaged packets.

In accordance with an alternative embodiment of the present invention, the window size and rate of change of window size are regulated by prioritization logic 401
25 rather than purely by TCP standards. For example, in an instance where TCP component 301 is maintaining ten TCP sockets 402, and each socket 402 is "entitled" to a maximum window size of ten under current bandwidth conditions, there are a total of 10x10 or 100 window
30 credits that can be managed. Conventional TCP would allow the TCP protocol to manage these independently. However, the present invention contemplates that credits may be removed from one TCP socket and given to another such that

a low priority socket may be given a maximum window size of 5 and a higher priority socket a maximum window size of 15. Moreover, the rate at which the window size increases/decreases under various conditions can be managed on a socket-by-socket basis such that in aggregate TCP unit 301 behaves as expected, but individually the TCP sockets 402 are micromanaged to provide prioritization of some sockets over others.

Controlling the rate and/or order that data is communicated with the last-mile communication system, as described above, is one method of controlling the prioritization and quality of service within the last-mile communication network that does not require any explicit control over all data between the Internet and the last-mile system itself. However, it is contemplated that in many instances the last-mile communication system provider will benefit from such external prioritization control and allow explicit management. Referring to FIG. 2, router 235 serves as the access point for a wireless interface component 215. Router 235 provides a control interface that enables external prioritization control provided by a front-end 201 to manipulate the rate and order at which router 235 handles packets. When such control is enabled, buffer resources within router 235 can be used to select which packets are buffered and which are sent to clients 205 according to the prioritization logic 401 shown in Fig. 4. It is not required, but useful if front-end 201 receives all data traffic as well. Alternatively, as shown in Fig. 2, the last-mile network can communicate in a conventional manner with the Internet at large.

As yet another alternative, a front-end 201 may be coupled between a router 235 and a wireless front-end 215 (or other shared bandwidth interface such as a cable

front-end 225). This requires specific cooperation of the last-mile communication provider, but enables a prioritization in a manner as described above.

While the present invention is primarily described in conjunction with client-server systems, the teachings are readily extended to peer-to-peer and server-to-server connections as well. With the increased availability of shared bandwidth connection channels it is increasingly feasible to couple both client, server, and peer processes through a shared bandwidth connection. Accordingly reference to "clients" and "servers" throughout the instant specification and claims are intended to be viewed broadly to include both peer-to-peer and server-to-server implementations.

Although the invention has been described and illustrated with a certain degree of particularity, it is understood that the present disclosure has been made only by way of example, and that numerous changes in the combination and arrangement of parts can be resorted to by those skilled in the art without departing from the spirit and scope of the invention, as hereinafter claimed. For example, while devices supporting HTTP data traffic are used in the examples, the HTTP devices may be replaced or augmented to support other public and proprietary protocols and languages including FTP, NNTP, SMTP, SQL and the like. In such implementations the front-end 201 and/or back end 203 are modified to implement the desired protocol. Moreover, front-end 201 and back-end 203 may support different protocols and languages such that the front-end 201 supports, for example, HTTP traffic with a client and the back-end supports a DBMS protocol such as SQL. Such implementations not only provide the advantages of the present invention, but also enable a client to

access a rich set of network resources with minimal client software.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237